IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants	:	Fraenkel, et al.
Appl. No.	:	10/057,295
Filed	:	October 19, 2001
For	:	POST-DEPLOYMENT MONITORING AND ANALYSIS OF SERVER PERFORMANCE
Examiner	:	Michael Young Won
Group Art Unit	:	2155

APPEAL BRIEF

United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

Appellants, Applicants in the above-captioned patent application, appeal the final rejection of Claims 1-39 set forth in the final Office Action mailed on February 15, 2006 (hereinafter "the Final Office Action"). If the fees paid via the EFS system are insufficient, please charge any additional fees that may be required to Deposit Account No. 11-1410.

I. REAL PARTY IN INTEREST

The real party in interest in the present application is Mercury Interactive Corporation.

II. RELATED APPEALS AND INTERFERENCES

No related appeals, interferences or judicial proceedings are currently pending.

III. STATUS OF CLAIMS

Claims 1-39, which are attached hereto as an appendix, are currently pending in the application and are the subject of this appeal.

IV. STATUS OF AMENDMENTS

No amendments have been made in response to the Final Office Action.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The present application includes five independent claims. Each independent claim is summarized below, with citations to corresponding portions of the specification and drawings as required by 37 C.F.R. § 41.37(c)(1)(v). These citations are provided to illustrate specific examples and embodiments of the recited claim language, and are not intended to limit the claims.

Independent Claim 1 is directed to a method of monitoring the operation of deployed web site system (30 in Figs. 1 and 26). The method involves monitoring the response times of the web site system (30) from multiple geographic locations, while monitoring server resource utilization parameters associated with the web site system (30) from a computer (166 in Fig. 26) that is local to the web site system (30). (See, e.g., page 6, lines 19-30; Fig. 26, and page 37, line 27 to page 41, line 13). The response times and server resource utilization parameters are automatically analyzed to evaluate whether a correlation exists between changes in the response times and changes in values of the server resource utilization parameters (see, e.g., page 7, lines 1-10; page 43, lines 11-25; and page 44, line 23 to page 46, line 15. See also Figs. 30, 36A and 36B for examples reports that may be generated based on the automated analysis).

Independent Claim 13 is directed to a system for monitoring performance of a deployed transactional server (30 in Fig. 1). The system comprises a first agent (32 in Figs. 1 and 26) and a second agent (166 in Fig. 26). The first agent (32) is configured to monitor the transactional server (30) over a network, and to collect performance data, including response times of the transactional server (see, e.g., page 11, lines 17-29 and page 38, lines 4-9). The second agent (166) is configured to monitor server resource utilization of the transactional server (30), and to collect data on one or more server resource utilization parameters (see, e.g., page 6, lines 19-30; and page 38, line 9 to page 40, line 26). The system also comprises an analysis component (168 in Fig. 26) that automatically detects correlations between response times of the transactional

server, as monitored by the first agent (32), and particular server resource utilization parameters, as monitored by the second agent (166). (See, e.g., page 7, lines 1-10; page 43, lines 11-25; and page 44, line 23 to page 46, line 15.)

Independent Claim 20 is directed to a method for monitoring the performance of a transactional server (30). The method comprises receiving performance data from a plurality of geographically distributed computers (40 in Figs. 1 and 26) that execute transactions on a transactional server (30) while monitoring associated response times (see, e.g., page 11, lines 17-29 and page 38, lines 4-9). The method also comprises receiving server resource utilization data from a computer (166) that monitors server resource utilization of the transactional server (30) during execution of the transactions by the plurality of computers (see, e.g., page 38, line 9 to page 40, line 26; and Fig. 26). The method further comprises automatically analyzing the performance data and the server resource utilization data to detect correlations between the performance of the transactional server and one or more particular server resource utilization parameters. (See, e.g., page 7, lines 1-10; page 43, lines 11-25; and page 44, line 23 to page 46, line 15.)

Independent Claim 25 is directed to a method of monitoring the operation of a deployed transactional server (30). The method comprising monitoring response times of the transactional server as seen from multiple geographic locations (see, e.g., page 11, lines 17-29 and page 38, lines 4-9), and, concurrently, monitoring a plurality of server resource utilization parameters associated with the transactional server (see, e.g., page 38, line 9 to page 40, line 26). The method further comprises programmatically evaluating whether a correlation exists between changes in the response times and changes in values of the plurality of server resource utilization parameters over time. (See, e.g., page 7, lines 1-10; page 43, lines 11-25; and page 44, line 23 to page 46, line 15.)

Independent Claim 33 is directed to a computer-implemented method of analyzing the performance of a server system (30). The method comprises monitoring a first performance parameter of the server system (30) over a period of time to generate a series of values of the first performance parameter, and monitoring a second performance parameter of the server system (30) over said period of time to generate a series of values of the second performance parameter

(see, e.g., page 38, lines 4-15; page 45, line 27 to page 46, line 15; and Fig. 26; note that neither performance parameter has to be a server resource utilization parameter). The method further comprises automatically analyzing the values of the first and second performance parameters to evaluate whether a correlation exists between the first performance parameter and the second performance parameter (see, e.g., page 7, lines 1-10; page 45, line 27 to page 46, line 15; page 51, line 14 to page 52, line 10; and page 53, lines 6-15).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

The following grounds for rejection are to be reviewed on appeal:

1. The rejection of Claims 25 and 29 as being anticipated by Sweet et al. (U.S. Pat. 6,519,714); and

2. The rejection of Claims 1-9, 11-16, 18-24, 26, 28 and 30-39 under 35 U.S.C. 103(a) as being unpatentable over the combination of Sweet et al. and Fletcher et al. (U.S. Pat. 6,321,264).

The remaining claims, dependent Claims 10, 17 and 27, were rejected over the combination of Sweet et al., Fletcher et al. and Booman et al. (U.S. Pat. 6,216,169). These three dependent claims stand or fall with the claims from which they depend.

Appellants reserve the right to disqualify one or more of the references as prior art in the future.

VII. ARGUMENT

For the reasons set forth below, Appellants respectfully submit that the anticipation and obviousness rejections are improper.

1. Anticipation rejection of Claims 25 and 29 over Sweet et al.

A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference. See *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631; 2 USPQ2d 1051, 1053 (Fed. Cir. 1987). As explained below, Sweet et al. does not anticipate either Claim 25 or Claim 29.

Claim 25

Claim 25 reads as follows:

> 25. A method of monitoring the operation of a deployed transactional server, the method comprising:
>
> (a) monitoring response times of the transactional server as seen from multiple geographic locations, including locations that are geographically remote from each other and from the transactional server;
>
> (b) concurrently with (a), monitoring a plurality of server resource utilization parameters associated with the transactional server; and
>
> (c) programmatically evaluating whether a correlation exists between changes in the response times and changes in values of the plurality of server resource utilization parameters over time.

The anticipation rejection of Claim 25 is improper because, among other reasons, Sweet et al. does not disclose the monitoring of response times of a transactional server "as seen from multiple geographic locations, including locations that are geographically remote from each other." In connection with this claim language, the Examiner points to col. 1, lines 55-62 and col. 8, lines 42-45 of Sweet et al. Neither these nor any other portion of Sweet et al., however, discloses the monitoring of response times "as seen from multiple geographic locations" as claimed.

The anticipation rejection of Claim 25 is also improper because Sweet et al. does not disclose "concurrently with (a), monitoring a plurality of server resource utilization parameters associated with the transactional server." In connection with this claim language, a response time of a server or server application is not a "server resource utilization parameter" as used in the claims. This is evident from the specification and claims of the present application, which consistently distinguish between response times and server resource utilization parameters. See, e.g., page 6, lines 19-30; Tables 2-4 on pages 41-42; and page 53, lines 16-28 (note that the disclosed server resource utilization parameters differ from response times in that they generally reveal the overall utilization levels of particular server resources). Thus, the application response time, application availability, and processing response time parameters disclosed in the cited portions of Sweet et al. do not satisfy the "plurality of server resource utilization parameters" language of the claim.

The anticipation rejection of Claim 25 is also improper because Sweet et al. does not disclose "programmatically evaluating whether a correlation exists between changes in the response times and changes in values of the plurality of server resource utilization parameters over time." In this regard, Sweet et al. does not disclose an evaluation of whether a correlation exists between changes in response times and changes in values of any other parameters. The portions of Sweet et al. cited by the Examiner, namely col. 1, lines 14-21; col. 2, lines 37-41; col. 5, lines 5-12; and col. 8, lines 8-12, do not suggest otherwise.

For the foregoing reasons, the anticipation rejection of Claim 25 is improper and should be withdrawn.

Claim 29

Because Claim 29 depends from Claim 25, the anticipation rejection of Claim 29 is improper for the reasons explained above.

2. Obviousness Rejection of Claims 1-9, 11-16, 18-24, 26, 28 and 30-39 over combination of Sweet et al. and Fletcher et al.

In rejecting claims under 35 U.S.C. § 103, the Examiner bears the initial burden of presenting a prima facie case of obviousness. See In re Rijckaert, 9 F.3d 1531, 1532, 28 USPQ2d 1955, 1956 (Fed. Cir. 1993). The Examiner may satisfy this burden only by showing some objective teaching in the prior art or that knowledge generally available to one of ordinary skill in the art. In re Fine, 837 F.2d 1071, 1074, 5 USPQ2d 1596, 1598 (Fed. Cir. 1988).

As set forth below, the Examiner has failed to satisfy this burden with respect to each of the rejected claims. By declining to present arguments with respect to some of the dependent claims, Appellants do not imply that the limitations added by such claims are disclosed or suggested by the references.

Claim 1

Claim 1 reads as follows:

> 1. A method of monitoring the operation of a deployed web site system, the method comprising:
>
> (a) monitoring response times of a web site system as seen from multiple geographic locations, including locations that are geographically remote from each other and from the web site system;

(b) concurrently with (a), monitoring a plurality of server resource utilization parameters associated with the web site system from a computer that is local to the web site system; and

(c) automatically analyzing the response times and server resource utilization parameters as monitored in (a) and (b) over a selected time period to evaluate whether a correlation exists between changes in the response times and changes in values of the plurality of server resource utilization parameters.

The rejection of Claim 1 is improper because, among other reasons, Sweet et al. and Fletcher et al. do not individually or collectively teach or suggest monitoring response times of a web site system "as seen from multiple geographic locations, including locations that are geographically remote from each other." In connection with this claim language, the Examiner points to col. 1, lines 55-62 and col. 8, lines 42-45 of Sweet et al. Neither these nor any other portion Sweet et al., however, teaches or suggests the monitoring of response times "as seen from multiple geographic locations" as claimed.

The rejection of Claim 1 is also improper because Sweet et al. and Fletcher et al. do not teach or suggest "concurrently with (a), monitoring a plurality of server resource utilization parameters associated with the web site system from a computer that is local to the web site system." In connection with this claim language, the Examiner points to various portions of Sweet et al, and particularly col. 1, lines 18-21; col. 2, lines 34-42; col. 3, lines 14-18; and col. 5, lines 8-12. Although some of these portions disclose the monitoring of particular performance parameters associated with a web site system, these parameters are not "server resource utilization parameters" as claimed.

The rejection of Claim 1 is also improper because Sweet et al. and Fletcher et al. do not teach or suggest "analyzing the response times and server resource utilization parameters as monitored in (a) and (b) over a selected time period to evaluate whether a correlation exists between changes in the response times and changes in values of the plurality of server resource utilization parameters." The Examiner contends that Sweet et al. teaches such an analysis, citing col. 1, lines 14-21; col. 1, lines 62-65; col. 4, lines 8-9; col. 5, lines 5-12; col. 7, lines 28-31 and 41-44; col. 8, lines 8-12, and col. 10, lines 42-44. Appellants respectfully disagree. Although some of these portions of Sweet et al. refer generally to the analysis of data collected from

executing synthetic transactions, nothing in Sweet et al. suggests an analysis that involves evaluating whether a correlation exists between changes in the response times and changes in values of any other parameters.

Fletcher et al. does not overcome the above-noted deficiencies in Sweet et al. The Examiner does not contend otherwise, but rather apparently relies on Fletcher et al. only in connection with the limitation "automatically" in subparagraph (c).

Because Sweet et al. and Fletcher et al. do not teach or suggest all of the limitations of Claim 1, the Examiner has not established prima facie obviousness, and the rejection of Claim 1 is improper.

Dependent Claims 2, 4-9 and 31

Because Claims 2, 4-9 and 31 depend from Claim 1, the obviousness rejections of Claims 2, 4-9 and 31 are improper for the reasons explained above for Claim 1.

Dependent Claim 3

Because Claim 3 depends from Claim 1, the obviousness rejection of Claim 3 is improper for the reasons explained above for Claim 1. In addition, the rejection of Claim 3 is improper because Sweet et al. and Fletcher et al. do not teach or suggest the limitations added by Claim 3, namely "wherein (a) comprises passively monitoring traffic resulting from actual web site users in at least some of the multiple geographic locations."

In connection with Claim 3, the Examiner points to a portion of Sweet et al. that describes how "Synthetic Transactions" are used to monitor an application. In particular, the Examiner relies on the following sentence, which appears at col. 2, lines 38-42: "In at least some cases, the evaluations are taken from the perspective of an end user (which may be a human being or a software program), so that the evaluations reflect the application response time and availability that would likely be experienced by the end user." When read in isolation, this sentence may appear to suggest that Sweet et al. uses the Synthetic Transactions to monitor traffic resulting from an actual user of the application. However, the paragraph which immediately follows this sentence makes clear that Sweet et al.'s system does not operate in such a manner, but rather actively simulates the actions of a user:

> As described in detail below, a Synthetic Transaction is a simulated application transaction. In at least some cases, it is advantageous if the simulated transaction

> is generated by an intelligent agent software system to evaluate a computer resource. For example, in a typical case concerning client/server technology on a network, a Synthetic Transactions software agent emulates a client side of a client/server transaction so that what happens on the network and on a server is the same, or nearly the same, as what would happen if the transaction were initiated by an actual end user on the client side. The agent may execute at a point on the network that is highly representative of where an actual end user would be situated, so that the path of associated communication across the network closely simulates the path of communication that is, or likely would be, associated with an actual end user.

In view of the foregoing, Appellants submit that the Examiner has not shown that Sweet et al. teaches the limitations of Claim 3.

Dependent Claim 11

Because Claim 11 depends from Claim 1, the obviousness rejection of Claim 11 is improper for the reasons explained above for Claim 1. In addition, the rejection of Claim 11 is improper because Sweet et al. and Fletcher et al. do not teach or suggest the limitations added by Claim 11, namely "applying a statistical algorithm to a sequence of response time measurements resulting from (a) to automatically detect a degradation in performance." The Examiner takes the position that the limitations of Claim 11 are inherent in Sweet et al.'s disclosure at col. 5, lines 5-8, which states that the information generated from executing a Synthetic Transaction is sent to a central server "for further processing and analysis." Appellants respectfully disagree.

"To establish inherency, the extrinsic evidence 'must make clear that the missing descriptive matter is necessarily present in the thing described in the reference, and that it would be so recognized by persons of ordinary skill." *In re Robertson*, 169 F.3d 743, 745, 49 USPQ2d 1949, 1950-51 (Fed. Cir. 1999) (citations omitted). "In relying upon the theory of inherency, the examiner must provide a basis in fact and/or technical reasoning to reasonably support the determination that the allegedly inherent characteristic necessarily flows from the teachings of the applied prior art." *Ex parte Levy,* 17 USPQ2d 1461, 1464 (Bd. Pat. App. & Inter. 1990) (emphasis in original).

In the present case, the Examiner has not presented any evidence or reasoning to suggest that the "further processing and analysis" referred to in Sweet et al. would necessarily involve "applying a statistical algorithm to a sequence of response time measurements." Indeed, the

analysis mentioned in Sweet et al. could simply involve the comparison of specific response time measurements to a response time threshold to determine whether an alarm condition is met. Thus, the Examiner has not shown that the limitations of Claim 11 are inherent in Sweet et al.

Dependent Claim 12

Because Claim 12 depends from Claim 11, the rejection of Claim 12 is improper for the reasons provided above for Claims 1 and 11. In addition, the rejection of Claim 12 is improper because Sweet et al. and Fletcher et al. do not teach or suggest "processing server resource utilization measurements resulting from (b) to identify at least one server resource parameter having a correlation with the degradation in performance." The cited portions of Sweet et al. (col. 1, lines 36-51 and col. 2, lines 8-11) do not teach the identification of a server resource parameter having a correlation with a detected degradation in performance.

Dependent Claim 30

Because Claim 30 depends from Claim 1, the obviousness rejection of Claim 30 is improper for the reasons explained above for Claim 1. In addition, the rejection of Claim 30 is improper because Sweet et al. and Fletcher et al. do not teach or suggest the limitations added by Claim 30, namely "in response to detecting in (c) a correlation between a response time degradation and a particular server resource utilization parameter, providing a visual representation of said correlation to a user." The portion of Sweet et al. (col. 8, lines 11-12) relied on by the Examiner simply does not disclose this feature.

Independent Claim 13

Independent Claim 13 reads as follows:

> 13. A system for monitoring performance of a deployed transactional server, the system comprising:
>
> a first agent configured to monitor a transactional server over a network, the first agent collecting performance data including response times of the transactional server;
>
> a second agent configured to monitor server resource utilization of the transactional server, the second agent collecting data on one or more server resource utilization parameters, wherein the second agent monitors server resource utilization over a time period in which the first agent monitors the transactional server; and
>
> an analysis component that automatically detects correlations between response times of the transactional server as monitored by the

first agent and particular server resource utilization parameters as monitored by the second agent.

The rejection of Claim 13 is improper because, among other reasons, Sweet et al. and Fletcher et al. do not individually or collectively teach or suggest "a second agent configured to monitor server resource utilization of the transactional server, the second agent collecting data on one or more server resource utilization parameters, wherein the second agent monitors server resource utilization over a time period in which the first agent monitors the transactional server." As explained above, the response times monitored in Sweet et al. are not "server resource utilization" parameters. Thus, the Examiner's reliance on Sweet et al. in connection with this portion of Claim 13 is misplaced.

The rejection of Claim 13 is also improper because Sweet et al. and Fletcher et al. do not individually or collectively teach or suggest "an analysis component that automatically detects correlations between response times of the transactional server as monitored by the first agent and particular server resource utilization parameters as monitored by the second agent." In connection with this claim language (with the exception of the term "automatically"), the Examiner points to various portions of Sweet et al. Nothing in Sweet et al., however, suggests an analysis component that detects correlations between response times and anything else, let alone particular server resource utilization parameters.

Fletcher et al. does not overcome these deficiencies in Sweet et al.

Because Sweet et al. and Fletcher et al. do not teach or suggest all of the limitations of Claim 13, the rejection of Claim 13 is improper.

Dependent Claims 14, 15, 18 and 19

Because Claims 14, 15, 18 and 19 depend from Claim 13, the obviousness rejections of Claims 14, 15, 18 and 19 are improper for the reasons explained above for Claim 13.

Dependent Claim 16

Because Claim 16 depends from Claim 13, the obviousness rejection of Claim 16 is improper for the reasons explained above for Claim 13. In addition, the rejection of Claim 16 is improper because Sweet et al. and Fletcher et al. do not teach or suggest the following limitations added by Claim 16: "wherein the first agent passively monitors traffic between client computers

and the transactional server to measure the response times." Sweet et al's method of using Synthetic Transactions to monitor the transactional server does not involve passively monitoring traffic between client computers and the transactional server to measure the response time. Thus, the Examiner's reliance on Sweet et al.'s Synthetic Transactions in connection with Claim 16 is misplaced.

Independent Claim 20

Claim 20 reads as follows:

> 20. A method for monitoring the performance of a transactional server, the method comprising:
>
> receiving performance data from a plurality of computers geographically distributed across a network, the plurality of computers executing transactions on a transactional server while monitoring associated response times;
>
> receiving server resource utilization data from a computer that monitors server resource utilization of the transactional server during execution of the transactions by the plurality of computers; and
>
> automatically analyzing the performance data and the server resource utilization data to detect correlations between the performance of the transactional server and one or more particular server resource utilization parameters.

The rejection of Claim 20 is improper because, among other reasons, Sweet et al. and Fletcher et al. do not teach or suggest "receiving performance data from a plurality of computers geographically distributed across a network, the plurality of computers executing transactions on a transactional server while monitoring associated response times." In connection with this claim language, there is nothing in Sweet et al. to suggest running the software agent described therein on computers that are geographically distributed. Thus, the Examiner's reliance on Sweet et al. in connection with this portion of Claim 20 is improper.

The rejection of Claim 20 is also improper because Sweet et al. and Fletcher et al. do not teach or suggest "receiving server resource utilization data from a computer that monitors server resource utilization of the transactional server during execution of the transactions by the plurality of computers." In connection with this claim language, the monitoring of response times and availability of a transactional server, as disclosed in Sweet et al., does not constitute monitoring of "server resource utilization" of the transactional server.

The rejection of Claim 20 is also improper because Sweet et al. and Fletcher et al. do not teach or suggest "automatically analyzing the performance data and the server resource utilization data to detect correlations between the performance of the transactional server and one or more particular server resource utilization parameters." In connection with this claim language (with the exception of the term "automatically"), the Examiner again points to various portions of Sweet et al., namely col. 1, lines 14-21; col. 1, lines 62-65; col. 5, lines 5-12; col. 8, lines 8-12, and col. 10, lines 42-44. Although some of these portions of Sweet et al. mention that the data collected by the software agents may be analyzed, there is nothing in Sweet et al. that suggests the type of analysis described in Claim 20. For example, Sweet et al. does not teach or suggest an analysis that involves detecting correlations between the performance of the transactional server and one or more particular parameters.

Fletcher et al. does not overcome the deficiencies noted above in Sweet et al.

Because Sweet et al. and Fletcher et al. do not teach or suggest all of the limitations of Claim 20, the rejection of Claim 20 is improper.

Dependent Claims 21 and 32

Because Claims 21 and 32 depend from Claim 20, the obviousness rejections of Claims 21 and 32 are improper for the reasons explained above for Claim 20.

Dependent Claim 22

Because Claim 22 depends from Claim 20, the obviousness rejection of Claim 22 is improper for the reasons explained above for Claim 20. In addition, the rejection of Claim 22 is improper because Sweet et al. and Fletcher et al. do not teach or suggest the limitations added by Claim 22, namely "wherein the server resource utilization data includes central process [sic, processing] unit (CPU) utilization data associated with the transactional server." The Examiner appears to take the position that the limitations of Claim 22 are satisfied by Sweet et al.'s disclosure of a "processing response time" at col. 3, lines 14-18. Appellants respectfully disagree. Sweet et al. defines "processing response time" as the "time spent in processing by the application." While Sweet et al.'s processing response time may be influenced by the server's current CPU utilization, it is not itself a measure of CPU utilization, and thus is not "CPU

utilization data." (Note that the term "CPU utilization" is used in the specification to refer to the "percent of time that [the] CPU is utilized." See page 41, Table 2.)

Dependent Claim 23

Because Claim 23 depends from Claim 20, the obviousness rejection of Claim 23 is improper for the reasons explained above for Claim 20. In addition, the rejection of Claim 23 is improper because Sweet et al. and Fletcher et al. do not teach or suggest the limitations added by Claim 23, namely "wherein the server resource utilization data includes memory allocation data." Appellants respectfully disagree with the Examiner's apparent position that the limitations of Claim 23 are taught or suggested by Sweet et al.'s brief mention, at col. 8, lines 30-33, of the possibility of evaluating "mass storage performance." In this regard, there is nothing in Sweet et al. to suggest that the briefly mentioned evaluation of mass storage performance would involve the use of memory allocation data—let alone the use of such data in the manner required by the combination of Claims 20 and 23.

Dependent Claim 24

Because Claim 24 depends from Claim 20, the obviousness rejection of Claim 24 is improper for the reasons explained above for Claim 20. In addition, the rejection of Claim 24 is improper because Sweet et al. and Fletcher et al. do not teach or suggest the limitations added by Claim 24, namely "wherein the server resource utilization data includes at least one of the following: hits per second data, requests queued data, current connections data, connection attempts data, or disk utilization data." The portion of Sweet et al. (col. 3, lines 33-36) relied on by the Examiner simply does not teach the use of any of these types of server resource utilization data—either generally or in the manner required by the combination of Claims 20 and 24.

Dependent Claims 26 and 28

Claims 26 and 28 depend from independent Claim 25, discussed above. The obviousness rejections of Claims 26 and 28 are improper at least because Sweet et al. and Fletcher et al. do not teach or suggest the limitations of Claim 25. For example, Sweet et al. and Fletcher et al. do not teach or suggest "programmatically evaluating whether a correlation exists between changes in the response times and changes in values of the plurality of server resource utilization parameters over time."

Independent Claim 33

Claim 33 reads as follows:

> 33. A computer-implemented method of analyzing the performance of a server system, the method comprising:
>
> monitoring a first performance parameter of the server system over a period of time to generate a series of values of the first performance parameter, wherein the server system responds to requests from clients during said period of time;
>
> monitoring a second performance parameter of the server system over said period of time to generate a series of values of the second performance parameter; and
>
> automatically analyzing the values of the first and second performance parameters to evaluate whether a correlation exists between the first performance parameter and the second performance parameter.

The rejection of Claim 33 is improper because Sweet et al. and Fletcher et al. do not teach or suggest, in the context of the other limitations of the claim, "automatically analyzing the values of the first and second performance parameters to evaluate whether a correlation exists between the first performance parameter and the second performance parameter."

In connection with this claim language (except for the term "automatically"), the Examiner again points to various portions of Sweet et al., namely col. 1, lines 14-21; col. 1, lines 62-65; col. 5, lines 5-12; col. 8, lines 8-12, and col. 10, lines 42-44. Although some of these portions of Sweet et al. mention that the data collected by Sweet et al.'s software agents may be analyzed, there is nothing in Sweet et al. that suggests evaluating whether a correlation exists between a first performance parameter and a second performance parameter.

Fletcher et al. does not overcome this deficiency in Sweet et al., and the Examiner does not contend otherwise.

Because Sweet et al. and Fletcher et al. do not teach or suggest all of the limitations of Claim 33, the rejection of Claim 33 is improper.

Dependent Claims 34 and 37-39

Because Claims 34 and 37-39 depend from Claim 33, the obviousness rejections of Claims 34 and 37-39 are improper for the reasons explained above for Claim 33.

Dependent Claim 35

Claim 35 depends from Claim 34, which depends from Claim 33. Thus, the obviousness rejection of Claim 35 is improper for the reasons explained above for Claim 33. In addition, the rejection of Claim 35 is improper because Sweet et al. and Fletcher et al. do not teach evaluating whether a correlation exists between a response time parameter and a server resource utilization parameter, as required by the combination of Claims 33, 34 and 35. As discussed above, Sweet et al. (cited in connection with Claim 35) does not disclose the monitoring of a server resource utilization parameter, and does not teach or suggest evaluating whether a correlation exists between two parameters.

Dependent Claim 36

Claim 36 depends from Claim 34, which depends from Claim 33. Thus, the obviousness rejection of Claim 36 is improper for the reasons explained above for Claim 33. In addition, the rejection of Claim 36 is improper because Sweet et al. and Fletcher et al. do not teach or suggest evaluating whether a correlation exists between a response time parameter and an operating system resource parameter, as required by the combination of Claims 33, 34 and 36. The portions of Sweet et al. (col. 2, lines 34-42 and col. 3, lines 14-18) cited by the Examiner in connection with Claim 36 do not disclose an operating system resource parameter.

VIII. CONCLUSION

For the reasons set forth above, Appellants respectfully submit that the rejections of Claims 1-39 are improper, and request that these rejections be reversed.

Respectfully submitted,

KNOBBE, MARTENS, OLSON & BEAR, LLP

Dated: 8-8-06

By: 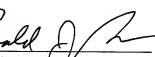

Ronald J. Schoenbaum
Reg. No. 38,297
2040 Main Street, 14th Floor
Irvine, CA 92614
949-721-2950

CLAIMS APPENDIX

1. A method of monitoring the operation of a deployed web site system, the method comprising:

(a) monitoring response times of a web site system as seen from multiple geographic locations, including locations that are geographically remote from each other and from the web site system;

(b) concurrently with (a), monitoring a plurality of server resource utilization parameters associated with the web site system from a computer that is local to the web site system; and

(c) automatically analyzing the response times and server resource utilization parameters as monitored in (a) and (b) over a selected time period to evaluate whether a correlation exists between changes in the response times and changes in values of the plurality of server resource utilization parameters.

2. The method of Claim 1, wherein (a) comprises monitoring the response times from agent computers in at least some of the multiple geographic locations.

3. The method of Claim 1, wherein (a) comprises passively monitoring traffic resulting from actual web site users in at least some of the multiple geographic locations.

4. The method of Claim 1, wherein (a) comprises generating page requests from a data center, and sending the page requests to the web site system via Internet points of presence located in at least some of the multiple geographic locations.

5. The method of Claim 1, wherein (b) comprises monitoring at least one server resource utilization parameter of a web server.

6. The method of Claim 1, wherein (b) comprises monitoring at least one server resource utilization parameter of an application server.

7. The method of Claim 1, wherein (b) comprises monitoring at least one server resource utilization parameter of a database server.

8. The method of Claim 1, wherein (b) comprises monitoring at least one server resource utilization parameter of a network device.

9. The method of Claim 8, wherein the network device is a router.

10. The method of Claim 8, wherein the network device is a bridge.

11. The method of Claim 1, further comprising applying a statistical algorithm to a sequence of response time measurements resulting from (a) to automatically detect a degradation in performance.

12. The method of Claim 11, further comprising processing server resource utilization measurements resulting from (b) to identify at least one server resource parameter having a correlation with the degradation in performance.

13. A system for monitoring performance of a deployed transactional server, the system comprising:

a first agent configured to monitor a transactional server over a network, the first agent collecting performance data including response times of the transactional server;

a second agent configured to monitor server resource utilization of the transactional server, the second agent collecting data on one or more server resource utilization parameters, wherein the second agent monitors server resource utilization over a time period in which the first agent monitors the transactional server; and

an analysis component that automatically detects correlations between response times of the transactional server as monitored by the first agent and particular server resource utilization parameters as monitored by the second agent.

14. The system of Claim 13, wherein the first agent is configured to monitor network hop delays.

15. The system of Claim 13, wherein the first agent sends request messages to the transactional server to measure the response times.

16. The system of Claim 13, wherein the first agent passively monitors traffic between client computers and the transactional server to measure the response times.

17. The system of Claim 13, further comprising a report generating component that generates reports associating the response times with the server resource utilization parameters by displaying the response times and the server resource utilization parameters on a time-

synchronized graph to permit a human operator to evaluate correlations detected by the analysis component.

18. The system of Claim 13, wherein the second agent is configured to monitor server resource utilization of a database server.

19. The system of Claim 13, further comprising an analysis component that automatically detects correlations between response times and server resource utilization parameters, wherein the analysis component analyzes sequences of values of said response times to automatically detect degradations in the performance of the transactional server.

20. A method for monitoring the performance of a transactional server, the method comprising:

receiving performance data from a plurality of computers geographically distributed across a network, the plurality of computers executing transactions on a transactional server while monitoring associated response times;

receiving server resource utilization data from a computer that monitors server resource utilization of the transactional server during execution of the transactions by the plurality of computers; and

automatically analyzing the performance data and the server resource utilization data to detect correlations between the performance of the transactional server and one or more particular server resource utilization parameters.

21. The method of Claim 20, wherein the performance data includes time stamps for associating the performance data and the server resource utilization data.

22. The method of Claim 20, wherein the server resource utilization data includes central process unit (CPU) utilization data associated with the transactional server.

23. The method of Claim 20, wherein the server resource utilization data includes memory allocation data.

24. The method of Claim 20, wherein the server resource utilization data includes at least one of the following: hits per second data, requests queued data, current connections data, connection attempts data, or disk utilization data.

25. A method of monitoring the operation of a deployed–transactional server, the method comprising:

(a) monitoring response times of the transactional server as seen from multiple geographic locations, including locations that are geographically remote from each other and from the transactional server;

(b) concurrently with (a), monitoring a plurality of server resource utilization parameters associated with the transactional server; and

(c) programmatically evaluating whether a correlation exists between changes in the response times and changes in values of the plurality of server resource utilization parameters over time.

26. The method of Claim 25, wherein (c) comprises automatically analyzing response time data and server resource utilization data resulting from (a) and (b), respectively.

27. The method of Claim 26, further comprising displaying, for a selected time window, a graph of the response times and a graph of at least one of the server resource utilization parameters.

28. The method of Claim 26, wherein (c) comprises analyzing response time data and server resource utilization data resulting from (a) and (b) with an automated analysis system that automatically detects correlations.

29. The method of Claim 25, wherein the transactional server is a web site system.

30. The method of Claim 1, further comprising, in response to detecting in (c) a correlation between a response time degradation and a particular server resource utilization parameter, providing a visual representation of said correlation to a user.

31. A computer system programmed to perform the method of Claim 1.

32. A computer system programmed to perform the method of Claim 20.

33. A computer-implemented method of analyzing the performance of a server system, the method comprising:

monitoring a first performance parameter of the server system over a period of time to generate a series of values of the first performance parameter, wherein the server system responds to requests from clients during said period of time;

monitoring a second performance parameter of the server system over said period of time to generate a series of values of the second performance parameter; and

automatically analyzing the values of the first and second performance parameters to evaluate whether a correlation exists between the first performance parameter and the second performance parameter.

34. The method of Claim 33, wherein the first performance parameter is a response time parameter.

35. The method of Claim 34, wherein the second performance parameter is a server resource utilization parameter.

36. The method of Claim 34, wherein the second performance parameter is an operating system resource parameter.

37. The method of Claim 33, wherein the step of automatically analyzing the values of the first and second performance parameters is performed in response to a user action.

38. A computer system programmed to perform the method of Claim 33.

39. A computer program which embodies the method of Claim 33 represented in computer storage.

EVIDENCE APPENDIX

None

RELATED PROCEEDINGS APPENDIX

None